

Predicting performance for post-quantum encrypted-file systems

Daniel J. Bernstein

University of Illinois at Chicago, USA
Ruhr University Bochum, Germany
Academia Sinica, Taiwan

Abstract Public-key cryptography is widely deployed for encrypting stored files. This paper uses microbenchmarks and purchase costs to predict the performance of various post-quantum KEMs in this application. In particular, this paper concludes that Classic McEliece is (1) the most efficient option and (2) easily affordable. As a quantitative example, the estimated five-year per-user cost of `mceliece6960119f` is 0.00024 dollars if 100000 files are encrypted and stored for an average user during those five years and each file is decrypted twice on average.

1 Introduction

Windows Pro, Enterprise, and Education include an “Encrypting File System”. Microsoft describes EFS in [Bra09] as “the built-in file encryption tool for Windows file systems”, and explains that each file stored under EFS is automatically encrypted to the public key of the user authorized to access the file:

EFS enables transparent encryption and decryption of files by using advanced, standard cryptographic algorithms. Any individual or program that doesn’t possess the appropriate cryptographic key cannot read the encrypted data. Encrypted files can be protected even from those who gain physical possession of the computer that the files reside on. . . . EFS encryption doesn’t occur at the application level but rather at the file-system level; therefore, the encryption and decryption process is transparent to the user and to the application. . . . File encryption uses a symmetric key, which is then itself encrypted with the public key of a public key encryption pair. The related private key must be available in order for the file to be decrypted. This key pair is bound to a user identity and made available to the user who has possession of the user ID and password.

Windows also supports a “BitLocker” mechanism for full-disk encryption. Microsoft’s BitLocker FAQ [Mic23] explains that per-file encryption and full-disk encryption provide different security features:

E-mail: djb@cr.yp.to (Daniel J. Bernstein)

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Excellence Strategy of the German Federal and State Governments—EXC 2092 CASA—390781972 “Cyber Security in the Age of Large-Scale Adversaries”; by the U.S. National Science Foundation under grant 2037867; and by the Taiwan’s Executive Yuan Data Safety and Talent Cultivation Project (AS-KPQ-109-DSTCP). “Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation” (or other organizations). Permanent ID of this document: 27c1d851ac2f907d715a0bbad5fa60623088bc54. Date: 2024.03.27.

Encrypting File System (EFS) can be used to encrypt files on a BitLocker-protected drive. BitLocker helps protect the entire operating system drive against offline attacks, whereas EFS can provide additional user-based file level encryption for security separation between multiple users of the same computer. EFS can also be used in Windows to encrypt files on other drives that aren't encrypted by BitLocker.

There are many other examples of per-file public-key encryption. Apple's iPhone, despite its reputation as a single-user device protected by full-disk encryption, also automatically creates ciphertexts to the user's Curve25519 public key, for example for backups and to be able to encrypt incoming messages while the device is locked. See [App22, pages 80 and 82]. For Linux tools that automatically encrypt files to a per-user GPG public key, see, e.g., [Car10], [Klo19], and [Phi19] for incoming mail messages, or [Ryg23] for incoming HTTP uploads.

To protect against an attacker with a quantum computer, one can and should upgrade encrypted-file systems from pre-quantum public-key encryption to post-quantum public-key encryption: for example, choose any post-quantum KEM, generate a KEM ciphertext to the user's KEM public key for each encrypted file, and use the resulting KEM session key to encrypt the file, or to encrypt another (perhaps longer-lasting) symmetric key used in turn to encrypt the file. "Hybrid"/"double encryption" approaches, retaining the existing layer of public-key encryption while adding a new layer of post-quantum public-key encryption, are also straightforwardly applicable here: e.g., encrypt with the old layer and then the new, or vice versa, or hash session keys and ciphertexts obtained from the old and new layers to obtain a hybrid session key used to encrypt the file.

1.1. Contributions of this paper. This paper asks what the performance consequences are of upgrading encrypted-file systems to use post-quantum cryptography. This paper uses microbenchmarks from eBACS [BerL24] to predict the answer for post-quantum KEMs currently under consideration by ISO (according to [Lam22, page 20]) and/or NIST. These KEMs are BIKE, Classic McEliece, FrodoKEM, HQC, and Kyber.

This comparison is analogous to, e.g., [DuHKL+23, abstract and Table 2 and Table 3] comparing some post-quantum KEMs according to "compactness" and "the round-trip time of ephemeral key exchange". The latter is computed in [DuHKL+23] as the sum of microbenchmarks for key generation, encapsulation, and decapsulation, disregarding overheads such as end-to-end communication latency. Two differences between this paper's comparison and the comparison in [DuHKL+23] are as follows:

- Beyond separately measuring storage space, communication, and CPU time, this paper uses purchase costs as a way to put all of these on the same scale, whereas [DuHKL+23] treats space and time as independent, not even accounting for the effect of "compactness" upon round-trip time.
- An encrypted-file system reuses a user's public key for all of the files encrypted to that user, whereas [DuHKL+23] considers a protocol that generates a new key for every ciphertext. All of the KEMs considered in this paper are designed for IND-CCA2 security, allowing keys to be reused for many ciphertexts.

These differences change which KEMs score best. Consider, in particular, Classic McEliece, the post-quantum encryption system with by far the strongest security track record (as quantified in [Ber23]). This system is often claimed to be non-competitive in performance because it has much larger per-key costs than alternatives: e.g., 1047319 bytes for a `mceliece6960119` public key, compared to 1568 bytes for a `kyber1024` public key. However, Classic McEliece has much smaller per-ciphertext costs than alternatives: e.g., 194 bytes for a `mceliece6960119` ciphertext, compared to 1568 bytes for a `kyber1024` ciphertext. Classic McEliece ends up being the most efficient way to encrypt many files to one key.

1.2. Variables. The predictions throughout this paper are in terms of the following variables. There are U users of the file system; e.g., $U = 1$ for a single-user file system. There are E files encrypted to an average user, for a total of UE encrypted files. There are D file-decryption operations by an average user, for a total of UD decrypted files. There are S files encrypted to an average user and still stored, for a total of US stored encrypted files. There is then a choice of KEM.

This paper takes four concrete examples to illustrate the impact of variations in E , D , and S :

- $(E, D, S) = (100000, 200000, 50000)$. This models a user receiving and storing 100 messages per work day for 5 years in an encrypted-file system, assuming 200 work days per year, while decrypting each message twice on average. S is half as large as E since at an average moment only half of the messages have arrived.
- $(E, D, S) = (200000, 400000, 100000)$. This models a more active user.
- $(E, D, S) = (200000, 200000, 100000)$. Same, except that the user decrypts each message only once on average.
- $(E, D, S) = (200000, 200000, 2000)$. Same, except that the user has a message-deletion policy and deletes an average message in just 10 work days. (Of course, this policy does not stop a malicious file-storage device from recording ciphertexts for longer. Such ciphertexts will remain decryptable with the user's secret key. Secure file erasure is outside the scope of this paper and does not appear to be a common feature of encrypted-file systems.)

This paper presumes that equipment is replaced every 5 years, and that new equipment will improve costs enough to outweigh increases in the number of messages, so this paper focuses on costs in the first 5 years.

Here is a sanity check on the 100-message-per-day number: [Sta23] indicates that there were 333.2 billion email messages sent worldwide each day in 2022, the equivalent of 600 billion messages per work day. Of course, email messages are not necessarily stored in encrypted-file systems (even if they should be for security); in the opposite direction, email messages are only one example of files that can be stored in encrypted-file systems. It would be interesting to collect real-world data regarding the variables (E, D, S) .

2 Space

This section predicts the space consumption of various KEMs in the context of encrypted-file systems. This section also reviews the purchase cost of storage.

2.1. Bytes stored. Each user incurs KEM-specific storage for a KEM public key and a KEM private key, so there are U public keys and U private keys. Each stored file incurs KEM-specific storage for a KEM ciphertext, so there are US ciphertexts. The KEM-specific storage consumed *per user* is thus 1 public key plus 1 private key plus S ciphertexts. Table 2.1.1 reports this amount of storage for the examples of S listed in Section 1.2; it is interesting to note the reversal of ordering between `mceliece8192128f` and `kyber512` as S changes. All tables in this paper are produced by Python scripts provided as PDF attachments to this paper.

2.2. The purchase cost of storage per byte. A typical 20TB hard drive today, the Western Digital WD201KFGX, is available for under \$400. The data sheet for this drive [Wes23a] indicates average power consumption of 6.9 watts during read/write and 1.6 watts on standby. If the drive runs continuously for its 5-year warranty period then it consumes between 70 and 300 kilowatt-hours depending on how frequent the read/write

Table 2.1.1: Number of bytes stored by each KEM per user of an encrypted-file system. S is the number of files stored per user. Rows are sorted by the last column. Columns are sorted by S .

	$S = 2000$	$S = 50000$	$S = 100000$
mceliece348864f	459612	5067612	9867612
mceliece460896f	849768	8337768	16137768
mceliece6960119f	1449267	10761267	20461267
mceliece6688128f	1474924	11458924	21858924
mceliece8192128f	1787944	11771944	22171944
kyber512	1538432	38402432	76802432
kyber768	2179584	54403584	108803584
kyber1024	3140736	78404736	156804736
bikel1	3152764	78656764	157306764
bikel3	6243188	155763188	311513188
hqc128round4	8870554	221654554	443304554
hqc192round4	17965108	448909108	897809108
frodokem640aes	19469504	486029504	972029504
frodokem640shake	19469504	486029504	972029504
hqc256round4	28856562	721064562	1442114562
frodokem976aes	31534928	787246928	1574446928
frodokem976shake	31534928	787246928	1574446928
frodokem1344aes	43328608	1081664608	2163264608
frodokem1344shake	43328608	1081664608	2163264608

operations are. A typical power supply is over 90% efficient, and then purchasing this energy at a typical price of 10 cents per kilowatt-hour adds between \$7 and \$33, making little difference next to the \$400 purchase price of the drive itself.

During the same 5-year period, the drive provides 100 terabyte-years of storage; i.e., one terabyte-year of storage costs about \$4. In other words, storing a byte for a year costs about 2^{-38} dollars.

Costs vary depending on the storage technology. For example, SSDs are often chosen over hard drives because they provide data access with lower latency and higher throughput; but a 4TB SSD currently costs about \$200, more than twice as expensive per byte as the hard drive. Also, a typical SSD warranty is limited to 5 years *or* a few petabytes written (see, e.g., [Wes23b]), so applications that rewrite all their data every day (e.g., continuously recording video feeds and retaining them for 24 hours) will end up replacing SSDs roughly every year, further increasing the cost per byte-year of storage.

Another reason for higher costs per byte-year is overprovisioning. For example, a user purchasing much more storage for a laptop than currently needed, so as to leave room for future growth, ends up paying more per byte-year of *available* storage. On the other hand, storing data in otherwise unused bytes has no effect on the purchase price of a disk if the same disk would have been purchased anyway. Precise accounting for the marginal cost of storing a byte would thus require investigating not just the extent of overprovisioning but also the interaction between the number of bytes stored and purchase costs. This paper focuses on the simpler scenario of storage being bought as necessary, for example by a company that aggregates storage across enough users to be able to predict usage.

Table 3.1.1: Number of bytes read+written by each KEM per user of an encrypted-file system. E is the number of files encrypted per user. D is the number of files decrypted per user. Rows are sorted by the last column. Columns are sorted by (E, D) .

	$E = 100000$ $D = 200000$	$E = 200000$ $D = 200000$	$E = 200000$ $D = 400000$
mceliece348864f	29067612	38667612	57867612
mceliece460896f	47337768	62937768	94137768
mceliece6960119f	59261267	78661267	117461267
mceliece6688128f	63458924	84258924	125858924
mceliece8192128f	63771944	84571944	126171944
kyber512	230402432	307202432	460802432
kyber768	326403584	435203584	652803584
kyber1024	470404736	627204736	940804736
bikel1	471906764	629206764	943806764
bikel3	934513188	1246013188	1869013188
hqc128round4	1329904554	1773204554	2659804554
hqc192round4	2693409108	3591209108	5386809108
frodokem640aes	2916029504	3888029504	5832029504
frodokem640shake	2916029504	3888029504	5832029504
hqc256round4	4326314562	5768414562	8652614562
frodokem976aes	4723246928	6297646928	9446446928
frodokem976shake	4723246928	6297646928	9446446928
frodokem1344aes	6489664608	8652864608	12979264608
frodokem1344shake	6489664608	8652864608	12979264608

3 Communication

This section predicts the amount of data read and written by various KEMs in the context of encrypted-file systems. This section also reviews the purchase cost of the relevant communication mechanisms.

3.1. Bytes written and read. Each user incurs the cost of writing a KEM public key and a KEM private key, so there are U public keys written and U private keys written. This paper assumes that caching of user keys in RAM is sufficiently effective that the cost of reading keys can be ignored.

Each encrypted file incurs the cost of writing a KEM ciphertext, so there are UE ciphertexts written. Each file-decryption operation incurs the cost of reading a KEM ciphertext, so there are UD ciphertexts read.

The KEM-specific reading and writing *per user* is thus 1 public key plus 1 private key plus $E + D$ ciphertexts. Table 3.1.1 computes this number of bytes for the examples of (E, D) listed in Section 1.2.

3.2. The purchase cost of communication per byte. The drive described in Section 2.2 has an internal transfer rate of 268 MB/second, according to [Wes23a]. Reading or writing data adds 5.3 watts compared to standby and, perhaps more relevant to marginal cost accounting, 3.1 watts compared to idle. Either way, the cost is roughly 2^{-51} dollars per byte at 10 cents per kWh.

Remote file systems are more expensive. For example, storing data in the cloud incurs costs to transmit data to and from the cloud, although it has the benefit of reducing the need for overprovisioning.

One way to estimate the purchase price of a byte of Internet communication is to look at advertisements of, e.g., \$70/month for a gigabit (roughly 100MB/second) connection, suggesting a cost around 2^{-42} dollars per byte. See, e.g., [IscW24]. On the other hand, [Gre20, page 20] says that “price discrimination based on usage of data, combined with data caps, is common in both wireline and wireless contracts”. A \$75/month “1,000 Mbps” connection limited to 1.2 terabytes per month (as in [Ger23]) would cost roughly 2^{-34} dollars per byte.

This paper uses an intermediate estimate obtained as follows. An estimate from 2018 [AslMKF18] was that Internet data transmission had an energy cost of “0.06 kWh/GB” and had a historical trend of dropping by a factor 2 every 2 years. At 10 cents per kWh, this is 0.006 dollars per GB in 2018, or projected 2^{-40} dollars per byte in 2024.

Recall from Section 2.2 the estimate of about 2^{-38} dollars per byte-year stored. If data is stored for only 10 days, about 2^{-5} years, then this storage costs 2^{-43} dollars per byte, so it is outweighed by 2^{-40} dollars per byte to upload to the cloud (plus the same for each download); communication costs for cloud storage end up dominant for, e.g., $(E, D, S) = (200000, 200000, 2000)$. If data is instead stored for years then a few uploads and downloads are outweighed by the data-storage costs: the data-storage costs end up dominant for, e.g., $(E, D, S) = (200000, 400000, 100000)$.

4 CPU time

This section predicts the number of CPU cycles used by various KEMs in the context of encrypted-file systems. This section also reviews the purchase cost of CPU cycles.

4.1. Cycles used. Each user generates a KEM public key and a KEM private key, so there are U key-generation operations. Each encrypted file incurs a KEM-specific encapsulation operation, so there are UE encapsulation operations. Each decryption of a file incurs a KEM-specific decapsulation operation, so there are UD decapsulation operations.

The KEM operations *per user* are thus 1 key generation plus E encapsulations plus D decapsulations. Table 4.1.1 shows the number of Skylake cycles for these operations for each (E, D) listed in Section 1.2, starting from the median cycle counts from [BerL24].

4.2. The purchase cost of CPU time per cycle. As in Section 2.2, this paper assumes that resources are purchased as needed. Checking <https://dell.com> on 7 February 2024 found a price of \$3616 for a Dell PowerEdge T440 server with two Intel Xeon Silver 4216 CPUs, 16GB of ECC RAM, and a 5-year warranty. This is not the lowest-cost processing currently available, but it suffices for purposes of this paper.

Each of the two CPUs has 16 cores and a base frequency of 2.1GHz; see [Int24]. This paper assumes that Turbo Boost is disabled. The CPU cores together carry out $2^{63.2}$ cycles in 5 years if they are used continuously. The CPU microarchitecture is Cascade Lake, which is essentially Skylake plus AVX-512; see [Wik22]. Beware that Table 4.1.1 does not use AVX-512 and is presumably overestimating costs on these CPUs.

The server has a 495W power supply, but can be reasonably estimated to draw under 300W from the wall since each CPU has a thermal design power of 100W. Running continuously for 5 years (43.83 kilo-hours) then costs under \$1315 at 10 cents per kWh.

Overall a CPU cycle on these CPUs costs about 2^{-51} dollars. Recall from Section 3.2 that reading a byte from a local hard drive also costs roughly 2^{-51} dollars, but reading a byte from remote storage costs roughly 2^{-40} dollars; and recall from Section 2.2 that a byte-year of storage costs about 2^{-38} dollars.

Table 4.1.1: Number of Skylake cycles used by each KEM per user of an encrypted-file system. E is the number of files encrypted per user. D is the number of files decrypted per user. Rows are sorted by the last column. Columns are sorted by (E, D) .

	$E = 100000$ $D = 200000$	$E = 200000$ $D = 200000$	$E = 200000$ $D = 400000$
kyber512	9201823150	12797223150	18403623150
kyber768	13756839708	19075439708	27513639708
kyber1024	19399454741	26792254741	38798854741
mceliece348864f	25796641148	28990241148	51564441148
mceliece460896f	55033306382	61966806382	109974606382
mceliece6960119f	65983868069	78974368069	131782568069
mceliece6688128f	69972100574	81871500574	139717300574
mceliece8192128f	72685016357	86899916357	145132516357
hqc128round4	100733187992	123417087992	201466287992
hqc192round4	210468898716	261015998716	420937598716
bikel1	323636986439	334370986439	647273386439
hqc256round4	412608412869	509574212869	825216412869
frodokem640aes	544672731817	729580831817	1089344031817
bikel3	996348703301	1021692403301	1992695803301
frodokem976aes	1048419156700	1414372756700	2096835356700
frodokem640shake	1250626473749	1669868073749	2501249073749
frodokem1344aes	1731268077705	2319174577705	3462531177705
frodokem976shake	2660059729718	3559710329718	5320111129718
frodokem1344shake	4593855509652	6142640609652	9187696409652

5 Total costs

This section predicts the total dollar costs per user of various KEMs in the context of encrypted-file systems. These predictions are split into two scenarios covered in Section 3.2: data being stored locally on a hard drive, and data being stored remotely in the cloud.

5.1. The local-storage scenario. Table 5.1.1 and Figure 5.1.2 make predictions for the local-storage scenario. The cost of 5 years of storing a byte is estimated as 2^{-36} dollars; the cost of reading or writing a byte is estimated as 2^{-51} dollars; the cost of a cycle is estimated as 2^{-51} dollars. These estimates come from Sections 2.2, 3.2, and 4.2.

The minimum per-user cost in Table 5.1.1 is from `mceliece348864f` for $(E, D, S) = (200000, 200000, 2000)$; recall that this is the setting of 200 files per day being encrypted to a user (and then decrypted) but files being deleted after 10 days. Here is a spot-check on this table entry:

- `mceliece348864f` has a 96-byte ciphertext for each of the 2000 stored files, occupying 192000 bytes of storage. `mceliece348864f` also has a 261120-byte public key and a 6492-byte private key, so the total storage is 459612 bytes. Multiplying by the estimated 2^{-36} dollars per byte gives 6.688 microdollars.
- `mceliece348864f` also writes 96-byte ciphertexts for 200000 encryption operations, reads 96-byte ciphertexts for 200000 decryption operations, and writes a public key and a private key, for a total of 38667612 bytes read and written. Multiplying by the estimated 2^{-51} dollars per byte gives 0.017 microdollars. (In Section 5.2, this 2^{-51} changes to 2^{-40} , increasing this cost to 35.168 microdollars.)
- The median cycle counts for `mceliece348864f` from [BerL24] are 28841148 cycles for keygen, 31936 cycles for enc, and 112871 cycles for dec, so 1 keygen, 200000 enc,

Table 5.1.1: Estimated dollar costs for each KEM per user of an encrypted-file system in the local-storage scenario. E is the number of files encrypted per user. D is the number of files decrypted per user. S is the number of files stored per user. Rows are sorted by the last column. Columns are sorted by (S, E, D) . See Figure 5.1.2 for a graph.

	$E = 200000$ $D = 200000$ $S = 2000$	$E = 100000$ $D = 200000$ $S = 50000$	$E = 200000$ $D = 200000$ $S = 100000$	$E = 200000$ $D = 400000$ $S = 100000$
mceliece348864f	0.00001958	0.00008521	0.00015648	0.00016652
mceliece460896f	0.00003991	0.00014579	0.00026238	0.00028372
mceliece6960119f	0.00005620	0.00018593	0.00033286	0.00035633
mceliece6688128f	0.00005786	0.00019785	0.00035448	0.00038019
mceliece8192128f	0.00006465	0.00020361	0.00036127	0.00038715
kyber512	0.00002821	0.00056302	0.00112344	0.00112600
kyber768	0.00004038	0.00079793	0.00159196	0.00159581
kyber1024	0.00005788	0.00114976	0.00229399	0.00229946
bike11	0.00019465	0.00128854	0.00243788	0.00257698
bike13	0.00054513	0.00270954	0.00498739	0.00541888
hqc128round4	0.00018468	0.00327082	0.00650653	0.00654158
hqc192round4	0.00037894	0.00662715	0.01318235	0.01325417
frodokem640aes	0.00060904	0.00731584	0.01447062	0.01463125
frodokem640shake	0.00102662	0.00762934	0.01488819	0.01525826
hqc256round4	0.00064878	0.01067803	0.02121439	0.02135584
frodokem976aes	0.00108980	0.01192364	0.02354212	0.02384660
frodokem976shake	0.00204252	0.01263935	0.02449484	0.02527802
frodokem1344aes	0.00166428	0.01651201	0.03251341	0.03302308
frodokem1344shake	0.00336224	0.01778326	0.03421137	0.03556556

and 200000 dec consume a total of 28.990 billion cycles. Multiplying 28.990 billion cycles by the estimated 2^{-51} dollars per cycle gives 12.874 microdollars.

The total matches the 19.58 microdollars (0.00001958) in Table 5.1.1.

In the same $(E, D, S) = (200000, 200000, 2000)$ setting, **kyber512** is 28.21 microdollars per user; **mceliece460896f** is 39.91 microdollars per user; **kyber768** is 40.38 microdollars per user; **mceliece6960119f** is 56.20 microdollars per user; **kyber1024** is 57.88 microdollars per user. Note that precise comparisons are uncertain since the numbers come from rough estimates; the reason for reporting extra digits is to support spot-checks.

For the other (E, D, S) settings in Table 5.1.1, the KEM costs are considerably higher and are dominated by ciphertext storage. This also makes Kyber strikingly less efficient than Classic McEliece.

5.2. The cloud-storage scenario. Table 5.2.1 and Figure 5.2.2 make predictions for the cloud-storage scenario, with reading or writing a byte estimated as 2^{-40} dollars. The cost of 5 years of storing a byte is estimated as 2^{-36} dollars; the cost of reading or writing a byte is estimated as 2^{-40} dollars (this is what is different from Section 5.1); the cost of a cycle is estimated as 2^{-51} dollars. These estimates come from Sections 2.2, 3.2, and 4.2.

Classic McEliece wins in each of the (E, D, S) settings in Table 5.2.1. The closest competition is for $(E, D, S) = (200000, 200000, 2000)$, where **mceliece348864f** costs 54.73 microdollars per user, **mceliece6960119f** costs 127.70 microdollars per user, **kyber512** costs 307.47 microdollars per user, and **kyber1024** costs 628.04 microdollars per user.




Figure 5.1.2: Estimated dollar costs for each KEM per user of an encrypted-file system in the local-storage scenario. E is the number of files encrypted per user. D is the number of files decrypted per user. S is the number of files stored per user. Columns are sorted by (S, E, D) . See Table 5.1.1 for a table.

6 Costs in context

This section compares the costs of post-quantum encrypted-file systems to other file-system costs and to costs of other applications of post-quantum cryptography.

6.1. Other file-system costs. The numbers listed in Tables 5.1.1 and 5.2.1 for `mceliece6960119f`, which is designed for long-term security, are under $1/2000$ of a dollar per user, suggesting that deployment in this application is very easily affordable. However, costs could be higher than this, for example because more expensive storage technologies are being used (see Section 2.2) or because many more files are being stored.

A different way to build confidence in the affordability of post-quantum encrypted-file systems is to compare the costs of post-quantum cryptography to other file-system costs. For example, measurements of the distribution of file sizes (see, e.g., [TanHB06]) and

Table 5.2.1: Estimated dollar costs for each KEM per user of an encrypted-file system in the cloud-storage scenario. E is the number of files encrypted per user. D is the number of files decrypted per user. S is the number of files stored per user. Rows are sorted by the last column. Columns are sorted by (S, E, D) . See Figure 5.2.2 for a graph.

	$E = 200000$ $D = 200000$ $S = 2000$	$E = 100000$ $D = 200000$ $S = 50000$	$E = 200000$ $D = 200000$ $S = 100000$	$E = 200000$ $D = 400000$ $S = 100000$
mceliece348864f	0.00005473	0.00011164	0.00019163	0.00021912
mceliece460896f	0.00009713	0.00018882	0.00031960	0.00036929
mceliece6960119f	0.00012770	0.00023980	0.00040436	0.00046310
mceliece6688128f	0.00013445	0.00025554	0.00043108	0.00049460
mceliece8192128f	0.00014153	0.00026158	0.00043815	0.00050185
kyber512	0.00030747	0.00077247	0.00140270	0.00154489
kyber768	0.00043600	0.00109465	0.00198759	0.00218924
kyber1024	0.00062804	0.00157738	0.00286415	0.00315470
bike11	0.00076663	0.00171753	0.00300987	0.00343495
bike13	0.00167782	0.00355906	0.00612008	0.00711791
hqc128round4	0.00179661	0.00447977	0.00811846	0.00895948
hqc192round4	0.00364353	0.00907560	0.01644694	0.01815105
frodokem640aes	0.00414346	0.00996666	0.01800503	0.01993286
frodokem640shake	0.00456103	0.01028016	0.01842260	0.02055987
hqc256round4	0.00589256	0.01461087	0.02645817	0.02922151
frodokem976aes	0.00681468	0.01621731	0.02926700	0.03243389
frodokem976shake	0.00776740	0.01693302	0.03021972	0.03386532
frodokem1344aes	0.00953017	0.02241144	0.04037930	0.04482189
frodokem1344shake	0.01122813	0.02368269	0.04207726	0.04736437

[DinN21]) indicate that the average file size has grown past a megabyte, making it implausible that there can be an issue with adding a relatively small KEM ciphertext to each file.

A counterargument is as follows. Both [TanHB06] and [DinN21] show a large spread of file sizes, with the average file size differing across use cases. For example, [DinN21] observes work by “IT staff (e.g., programmer, systems administrator)” storing much smaller files than other use cases, and explains this via a previous study of the prevalence of “plain text files”, “documents”, and “media” across different use cases.

It would be interesting to collect statistics on the sizes of files stored on encrypted-file systems: for example, how do incoming messages automatically encrypted by an iPhone compare in size to other files? It would also be interesting to collect statistics on how often files of different sizes are decrypted.

It is conceivable that there are use cases that store a huge number of 100-byte files in an encrypted-file system and that would object to adding 194 bytes per file for a mceliece6960119f ciphertext. Quantitatively, a dollar pays for 5-year storage of about a quarter billion of these ciphertexts; what happens if a user is storing many billions of files?

However, it is reasonably clear that such use cases are not common, since they would already incur larger overheads for unencrypted-file systems. For example, Windows NTFS allocates storage in units of “clusters”, and NTFS “clusters” cannot be smaller than 512 bytes (see [Wat19]), so every 100-byte file is wasting 412 bytes of disk space.

To summarize, the evidence indicates that upgrading encrypted-file systems to mceliece6960119f will be easily affordable for practically all use cases.




Figure 5.2.2: Estimated dollar costs for each KEM per user of an encrypted-file system in the cloud-storage scenario. E is the number of files encrypted per user. D is the number of files decrypted per user. S is the number of files stored per user. Columns are sorted by (S, E, D) . See Table 5.2.1 for a table.

6.2. Other applications. The affordability of this upgrade to post-quantum cryptography sounds very different from a common narrative that post-quantum performance is a major issue. Consider, for example, NIST’s round-3 report [AlaApCD+22] commenting extensively on post-quantum performance, deferring a decision on whether to standardize Classic McEliece, and rejecting FrodoKEM outright:

Classic McEliece was a finalist, but is not being standardized by NIST at this time. Although it is widely regarded as secure, NIST does not yet anticipate it being widely used because of its large public key size . . . Frodo is clearly not an immediate drop-in general-purpose scheme.

One might try to explain the different conclusions regarding affordability by hypothesizing that

- encrypted-file systems are an unusual application because of their reuse of public keys, while
- *normal* applications generate a new public key for every ciphertext, involving thousands of times more data for Classic McEliece than just the ciphertexts.

But it is very easy to find literature on protocols relying on IND-CCA2 security to be able to reuse keys. Furthermore, long-term identity keys are used pervasively in protocol design, and can be long-term signature keys but—in any protocol where the signer is online—can also be long-term encryption keys, as in, e.g., [BelCK98, Section 3.2], [Ber09], [Ber16], [Ber18, Section 8], and [SchSW20]. The available evidence indicates that post-quantum encryption is more efficient than post-quantum signatures. A full analysis of the efficiency of post-quantum cryptography certainly has to include long-term post-quantum encryption keys, contrary to the idea that each encryption key is used just once.

Even in situations where there *is* a new public key for every ciphertext, this merely doubles the amount of FrodoKEM data. Why is FrodoKEM “clearly not an immediate drop-in general-purpose scheme”? One might try to answer this question by formulating the following further hypotheses:

- Encrypted-file systems, at least for the limited sizes of (E, D, S) in Tables 5.1.1 and 5.2.1, are far below the overall volume of “general-purpose” usage of public-key cryptography, in particular for retrieving web pages.
- Quantitatively, even though those tables estimate 5 years of active use of an encrypted-file system as costing just pennies per user with FrodoKEM, users incur many more public-key operations for retrieving web pages, and with FrodoKEM or Classic McEliece those operations would add up to a *significant* cost.

But there does not appear to be any documentation quantifying the dollar cost here.

An earlier NIST report [AlaAIAC+20] stated that FrodoKEM in TLS key exchange would cost “around 20,000 bytes” plus “2 million cycles” for the server, and concluded that FrodoKEM does not have “acceptable performance in widely used applications overall”. In response to the question of why this cost was not “acceptable”, NIST wrote the following [Per20]:

While it is not possible to speak for what every user of our standards would or wouldn’t find “acceptable”, there is a pretty large difference between the performance of Frodo on the one hand and Kyber, NTRU, and Saber on the other hand. We are therefore more confident that Kyber, NTRU, or Saber will be considered “acceptable” for most users than that Frodo will.

Quantifying costs in context produces different conclusions. The average web page has grown past 2 megabytes (see [Htt23]), and a single TLS session can retrieve any number of web pages from a server. A web page can collect data from multiple servers, but the number of TCP connections per web page has dropped to 10 (see again [Htt23]), so the number of TLS sessions involved is at most 10. In other words, the 20000 bytes for FrodoKEM cannot add more than 10% to overall web-page traffic, and if a user is retrieving at least 10 web pages per TLS session then 20000 bytes cannot add more than 1%.

NIST’s calculation of 20000 bytes implicitly assumes that a key is transmitted for every ciphertext—which is another indication of the application not caring about key-distribution costs, given that those costs can obviously be reduced:

- A server that announces a new key every day can receive any number of ciphertexts to that key from any number of clients. Each client receives at most one key from the server per day—for example, spending at most 2^{-20} dollars per server per day for Classic McEliece keys.

- Keys also do not need to be sent from end to end: they can use lower-cost broadcast networks, for example using existing ISP-level caches to share key-retrieval costs across users. This can turn Classic McEliece into the lowest-cost KEM (see [AlbBCC+22, Section 4] and [Ber23, pages 16–18]), much as in the case of encrypted-file systems.
- Users are often on more expensive cellular networks, but they can download keys for frequently used servers in advance. An iPhone already automatically waits until it is on a Wi-Fi network before downloading software updates (see [App24]); it can do the same to reduce the cost of downloading new public keys. Servers can also distribute multiple keys in advance to simplify download scheduling, for example distributing the keys to be used for the next 7 days; this is compatible with erasing each key within a day of user data being sent to that key.

The absence of these features in TLS today is easily explained by (1) the history of TLS using small-key cryptosystems and (2) the lack of evidence that the key-distribution costs matter to the end user.

Recall the estimate from Section 3.2 of 2^{-40} dollars per byte for Internet communication. This corresponds to 50 million 20000-byte FrodoKEM sessions per dollar: i.e., spending a dollar on FrodoKEM traffic would pay for 50000 new FrodoKEM sessions per work day for 5 years. These are sizes for `frodokem640`, but `frodokem1344` is only about twice as large. It would be interesting to collect data on the number of TLS sessions initiated in a day by an average user’s browser, the number of different TLS servers contacted, etc.

As one data point, [Ker23] reported in 2023 that Cloudflare was handling 45 million HTTP connections per second, i.e., $2^{50.3}$ HTTP connections per year. If each of these had immediately added 20000 bytes for `frodokem640`—with no caching or reuse of keys—then the total cost of the added communication would have been roughly \$25 million/year. To put this in perspective, [Clo23] reported that Cloudflare’s “GAAP gross profit was \$742.6 million” for 2022.

In general, given trends towards larger volumes of data, one would expect that the costs of post-quantum public-key encryption will ultimately be irrelevant compared to the costs of symmetric cryptography and non-cryptographic costs. This is not necessarily the situation today, but it is concerning to see unsubstantiated claims of unaffordability being used to discourage deployment of lower-risk cryptosystems.

References

- [@21] — (no editor), *Information: equity, diversity, inclusion, justice, and relevance—proceedings of the 84th ASIS&T annual meeting, ASIST 2021, Salt Lake City, UT, USA, October 30–November 2, 2021*, Proceedings of the Association for Information Science and Technology, 58, Wiley, 2021. See [18].
- [AlaAlAC+20] Gorjan Alagic, Jacob Alperin-Sheriff, Daniel Apon, David Cooper, Quynh Dang, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, *Status report on the second round of the NIST Post-Quantum Cryptography Standardization Process* (2020), NISTIR 8309. URL: <https://csrc.nist.gov/publications/detail/nistir/8309/final>. Citations in this document: §6.2.
- [AlaApCD+22] Gorjan Alagic, Daniel Apon, David Cooper, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel

- Smith-Tone, *Status report on the third round of the NIST Post-Quantum Cryptography Standardization Process* (2022), NISTIR 8413. URL: <https://web.archive.org/web/20230824124130/https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8413-upd1.pdf>. Citations in this document: §6.2.
- [AlbBCC+22] Martin R. Albrecht, Daniel J. Bernstein, Tung Chou, Carlos Cid, Jan Gilcher, Tanja Lange, Varun Maram, Ingo von Maurich, Rafael Misoczki, Ruben Niederhagen, Kenneth G. Paterson, Edoardo Persichetti, Christiane Peters, Peter Schwabe, Nicolas Sendrier, Jakub Szefer, Cen Jung Tjhai, Martin Tomlinson, Wen Wang, *Classic McEliece: conservative code-based cryptography: guide for implementors* (2022). URL: <https://classic.mceliece.org/mceliece-impl-20221023.pdf>. Citations in this document: §6.2.
- [App22] Apple, *Apple Platform Security* (2022). URL: https://web.archive.org/web/20231128180844/https://help.apple.com/pdf/security/en_US/apple-platform-security-guide.pdf. Citations in this document: §1.
- [App24] Apple, *Update iOS on iPhone* (2024). URL: <https://web.archive.org/web/20240319105030/https://support.apple.com/guide/iphone/update-ios-iph3e504502/ios>. Citations in this document: §6.2.
- [AslMKF18] Joshua Aslan, Kieren Mayers, Jonathan G. Koomey, Chris France, *Electricity intensity of Internet data transmission: untangling the estimates*, *Journal of Industrial Ecology* **22** (2018), 785–798. URL: <https://onlinelibrary.wiley.com/doi/10.1111/jiec.12630>. Citations in this document: §3.2.
- [BelCK98] Mihir Bellare, Ran Canetti, Hugo Krawczyk, *A modular approach to the design and analysis of authentication and key exchange protocols (extended abstract)*, in *STOC 1998 [Vit98]* (1998), 419–428. URL: <https://eprint.iacr.org/1998/009>. Citations in this document: §6.2.
- [Ber09] Daniel J. Bernstein, *DNSCurve: Usable security for DNS* (2009). URL: <https://dnscurve.org>. Citations in this document: §6.2.
- [Ber16] Daniel J. Bernstein, *The post-quantum Internet* (2016). URL: <https://cr.yp.to/talks.html#2016.02.24>. Citations in this document: §6.2.
- [Ber18] Daniel J. Bernstein, *D2.5: Internet: Integration* (2018). URL: <https://www.pqcrypto.eu/deliverables/d2.5.pdf>. Citations in this document: §6.2.
- [Ber23] Daniel J. Bernstein, *Migrating to the McEliece cryptosystem* (2023). URL: <https://cr.yp.to/talks.html#2023.10.25>. Citations in this document: §1.1, §6.2.
- [BerL24] Daniel J. Bernstein, Tanja Lange (editors), *eBACS: ECRYPT Benchmarking of Cryptographic Systems* (2024), Accessed 4 March 2024. URL: <https://bench.cr.yp.to>. Citations in this document: §1.1, §4.1, §5.1.
- [BolK23] Alexandra Boldyreva, Vladimir Kolesnikov (editors), *Public-key cryptography—PKC 2023—26th IACR international conference on practice and theory of public-key cryptography, Atlanta, GA, USA,*

- May 7–10, 2023, proceedings, part I*, 13940, Springer, 2023. ISBN 978-3-031-31367-7. See [19].
- [Bra09] Roberta Bragg, *The Encrypting File System* (2009). URL: [https://web.archive.org/web/20221007214414/https://learn.microsoft.com/en-us/previous-versions/tn-archive/cc700811\(v=technet.10\)?redirectedfrom=MSDN#XSLTsection125121120120](https://web.archive.org/web/20221007214414/https://learn.microsoft.com/en-us/previous-versions/tn-archive/cc700811(v=technet.10)?redirectedfrom=MSDN#XSLTsection125121120120). Citations in this document: §1.
- [Car10] Mike Cardwell, *Automatically encrypting all incoming email* (2010). URL: https://www.grepular.com/Automatically_Encrypting_all_Incoming_Email. Citations in this document: §1.
- [Clo23] Cloudflare, *Cloudflare announces fourth quarter and fiscal year 2022 financial results* (2023). URL: <https://web.archive.org/web/20231208094317/https://www.cloudflare.net/news/news-details/2023/Cloudflare-Announces-Fourth-Quarter-and-Fiscal-Year-2022-Financial-Results/default.aspx>. Citations in this document: §6.2.
- [DinN21] Jesse David Dinneen, Ba Xuan Nguyen, *How big are peoples’ computer files? File size distributions among user-managed collections*, in ASIST 2021 [21] (2021), 425–429. URL: <https://arxiv.org/abs/2107.03272>. Citations in this document: §6.1, §6.1, §6.1.
- [DuHKL+23] Julien Duman, Kathrin Hövelmanns, Eike Kiltz, Vadim Lyubashevsky, Gregor Seiler, Dominique Unruh, *A thorough treatment of highly-efficient NTRU instantiations*, in PKC 2023 [BK23] (2023), 65–94. URL: <https://eprint.iacr.org/2021/1352>. Citations in this document: §1.1, §1.1, §1.1, §1.1, §1.1.
- [Ger23] Nick Gerhardt, *Xfinity Internet review: plans, prices and speed 2023* (2023). URL: <https://web.archive.org/web/20231210121119/https://www.forbes.com/home-improvement/internet/xfinity-internet-review/>. Citations in this document: §3.2.
- [Gre20] Shane Greenstein, *The basic economics of Internet infrastructure*, *Journal of Economic Perspectives* **34** (2020), 192–214. Citations in this document: §3.2.
- [Htt23] HTTP Archive, *Report: state of the web* (2023), accessed 1 December 2023. URL: <https://httparchive.org/reports/state-of-the-web>. Citations in this document: §6.2, §6.2.
- [Int24] Intel, *Intel Xeon Silver 4216 Processor* (2024). URL: <https://web.archive.org/web/20240317082918/https://ark.intel.com/content/www/us/en/ark/products/193394/intel-xeon-silver-4216-processor-22m-cache-2-10-ghz.html>. Citations in this document: §4.2.
- [IscW24] Lisa Iscrupé, Hannah Whatley, *What is 1 gigabit internet? Is it worth the cost?* (2024). URL: <https://web.archive.org/web/20240301102137/https://www.usatoday.com/tech/internet/should-you-purchase-internet-over-1-gig/>. Citations in this document: §3.2.

- [Ker23] Sean Michael Kerner, *Attackers target Microsoft Exchange, according to Cloudflare application security report* (2023). URL: <https://web.archive.org/web/20230316083215/https://www.sdxcentral.com/articles/news/attackers-target-microsoft-exchange-according-to-cloudflare-application-security-report/2023/03/>. Citations in this document: §6.2.
- [Klo19] Julian Andres Klode, *Encrypted email storage, or DIY ProtonMail* (2019). URL: <https://blog.jak-linux.org/2019/06/13/encrypted-email-storage/>. Citations in this document: §1.
- [Lam22] Brian LaMacchia, *The long road ahead to transition to post-quantum cryptography* (2022). URL: <https://web.archive.org/web/20230401090854/https://secdev.ieee.org/wp-content/uploads/2022/10/LaMacchia-Keynote-IEEESecDev2022.pdf>. Citations in this document: §1.1.
- [LigOKV20] Jay Ligatti, Xinming Ou, Jonathan Katz, Giovanni Vigna (editors), *CCS '20: 2020 ACM SIGSAC conference on computer and communications security, virtual event, USA, November 9–13, 2020*, ACM, 2020. ISBN 978-1-4503-7089-9. See [33].
- [Mic23] Microsoft, *BitLocker FAQ* (2023). URL: <https://web.archive.org/web/20230813064046/https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/faq>. Citations in this document: §1.
- [Per20] Ray Perlner, *Re: ROUND 2 OFFICIAL COMMENT: Frodo* (2020). URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/7aenKgDwV2k/m/1SJkflT9DQAJ>. Citations in this document: §6.2.
- [Phi19] J. Nathanael Philipp, *Auto encrypt all incoming email with postfix* (2019). URL: <https://jnphilipp.org/posts/auto-encrypt-all-incoming-email-with-postfix/>. Citations in this document: §1.
- [Ryg23] Leif Ryge, *Yet another solution to XKCD #949* (2023). URL: <https://github.com/leif/dropsite>. Citations in this document: §1.
- [SchSW20] Peter Schwabe, Douglas Stebila, Thom Wiggers, *Post-quantum TLS without handshake signatures*, in *CCS 2020 [LigOKV20]* (2020), 1461–1480. URL: <https://eprint.iacr.org/2020/534>. Citations in this document: §6.2.
- [Sta23] Statista, *Number of sent and received e-mails per day worldwide from 2017 to 2026* (2023). URL: <https://web.archive.org/web/20231127073251/https://www.statista.com/statistics/456500/daily-number-of-e-mails-worldwide/>. Citations in this document: §1.2.
- [TanHB06] Andrew S. Tanenbaum, Jorrit N. Herder, Herbert Bos, *File size distribution on UNIX systems: then and now*, *ACM SIGOPS Operating Systems Review* **40** (2006), 100–104. URL: <https://www.minix3.org/docs/jorrit-herder/osr-jan06.pdf>. Citations in this document: §6.1, §6.1.

- [Vit98] Jeffrey Scott Vitter (editor), *Proceedings of the thirtieth annual ACM symposium on the theory of computing, Dallas, Texas, USA, May 23–26, 1998*, ACM, 1998. ISBN 0-89791-962-9. See [8].
- [Wat19] Garrett Watumull, *Cluster size recommendations for ReFS and NTFS* (2019). URL: <https://techcommunity.microsoft.com/t5/storage-at-microsoft/cluster-size-recommendations-for-refs-and-ntfs/ba-p/425960>. Citations in this document: §6.1.
- [Wes23a] Western Digital, *WD Red Pro* (2023). URL: https://web.archive.org/web/20231026192501/https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/internal-drives/wd-red-pro-hdd/product-brief-western-digital-wd-red-pro-hdd.pdf. Citations in this document: §2.2, §3.2.
- [Wes23b] Western Digital, *WD internal SSD endurance and warranty periods* (2023). URL: https://web.archive.org/web/20231031195024/https://support-en.wd.com/app/answers/detailweb/a_id/30797/. Citations in this document: §2.2.
- [Wik22] WikiChip, *Cascade Lake—Microarchitectures—Intel* (2022). URL: https://en.wikichip.org/wiki/intel/microarchitectures/cascade_lake. Citations in this document: §4.2.